Embedded Intel® Architecture in the Carrier-Class Voice over IP Gateway

Krishna Shetty
Applications Architect
Embedded Intel Architecture Division
Intel Corporation

Rishi Mukhopadhyay
Engineering Project Manager
Embedded Intel Architecture Division
Intel Corporation

Embedded Intel® Architecture in the Carrier-Class Voice over IP Gateway

Contents

Executive Summary	3
Intel® Internet Exchange Architecture and Embedded Intel® Architecture in a Voice over IP (VoIP) Gateway	4
Voice over IP: General Considerations	5
Co-existence of VoIP and PSTN	5
Signaling in PSTN and IP	5
VoIP Protocols	6
H.323 Protocol and Revisions	6
Bridging Between PSTN and IP via the H.323 Gateway	6
Building a VoIP Gateway	7
PC vs. Embedded Solutions	7
Hardware Recommendations	8
Software Requirements	8
System Architecture	9
PCI Bandwidth Considerations for H.323 Signaling	10
Applied Computing Platform Providers and Software Suppliers	11
Conclusion	11
For More Information	11

Executive Summary

The convergence of voice and data is increasing the value and usefulness of the Internet in every walk of life. The combination of the PSTN (Public Switched Telephone Network) and the Internet can provide access to the massive information base within the Internet through a simple mechanism like a telephone handset. Such a revolution has already begun with the introduction of standard services for voice, fax, and data that are currently delivered by the PSTN over the IP-based packet networks. The promise of transporting traditional telephony services over IP (Internet Protocol) based networks is leading to the development of cost-effective gateway equipment based on embedded systems.

Voice over IP (VoIP) is essentially a way to carry voice using the IP as a transport mechanism, similar to the way data is carried in the network with TCP/IP. With the growing volume of data traffic, voice becomes yet another application over the data network.

VoIP currently focuses on two key applications. First is the corporate network that interconnects remote branch offices with data services and would like to add voice and fax services. Corporations are interested in VoIP solutions due to the cost savings that can result through unification of two disparate networks. Apart from reducing operational and network management costs, corporations can reduce the communication access charges and settlement fees involving multi-international sites.

The second application is VoIP over public networks. ISPs (Internet Service Providers) are increasingly interested in becoming ITSPs (Internet Telephone Service Providers). This would enable them to add voice and multi-media services, and break out from the current monthly fee structure for data-only services. In addition, emerging Competitive Local Exchange Carriers (CLECs) are building new-generation IP network facilities. These carriers are interested in VoIP because it is more efficient than the conventional voice networks. Finally, current POTS (Plain Old Telephone Service) providers are interested in VoIP in order to remain competitive and generate additional revenues from new services, including “follow me,” unified messaging and text-to-speech conversion.

In the short term, cost considerations may be the biggest motivator pushing VoIP adoption. In the long term, VoIP has potential for much more than just telephony. It promotes the convergence of services by offering flexible bandwidth for voice, fax, video, data sharing, white-boarding, and easy connectivity to various back-end services. VoIP will enable the convergence of infrastructure by using the network bandwidth to efficiently utilize coding, compression, and packet switching.

Intel® Internet Exchange Architecture and Embedded Intel® Architecture in a VoIP Gateway

As the demand for fast deployment of new services (such as VoIP services) and applications grow, fixed-function ASIC-based designs will fail to provide the needed flexibility and time-to-market advantages required by customers deploying new services. Next generation packet processing hardware will need to be both fast and easily programmed with a consistent software architecture supported by a rich assortment of hardware components offering a large spectrum of price/performance points. Intel® Internet Exchange Architecture (IXA) provides a new level of integration, performance and programmability for new applications, such as VoIP.

Intel Internet Exchange Architecture is an end-to-end family of high-performance, flexible and scalable hardware and software development building blocks designed to meet the growing performance requirements of today's networks. Based on programmable silicon and software building blocks, Intel IXA solutions enable faster development, more cost-effective deployment and future upgradability of network and communications systems.

The Intel IXA separates the control and service functions from the wire-speed packet forwarding aspects of the communication infrastructure. This demarcation is useful in maintaining the scalability of communication systems, while allowing the individual layers to be developed independently. Such architecture can accommodate new and enhanced services and features as they emerge independent of underlying technological advances in the forwarding layer, as seen in Figure 1.

Embedded processors play an important role in Intel IXA by meeting processing requirements in the control and services layer. In the case of a VoIP gateway, an embedded Intel® Architecture (EIA) processor can be used to handle the signaling stacks of both the PSTN and IP networks. In addition, EIA processors can run several other applications, including network management, supervisory, billing, and call record functions.

The I/O peripheral boards shown in Figure 1 are basically responsible for packet forwarding and other wire-speed communication functions that are closely coupled to the physical layer. In the case of a gateway between PSTN and IP networks, the peripheral boards can also handle the T1/E1 framing, ATM related functions, or other higher-speed PSTN interfaces. These I/O boards also provide the interface to the IP network and media related algorithms, RTP/RTCP protocols, and packetization. Depending on the port

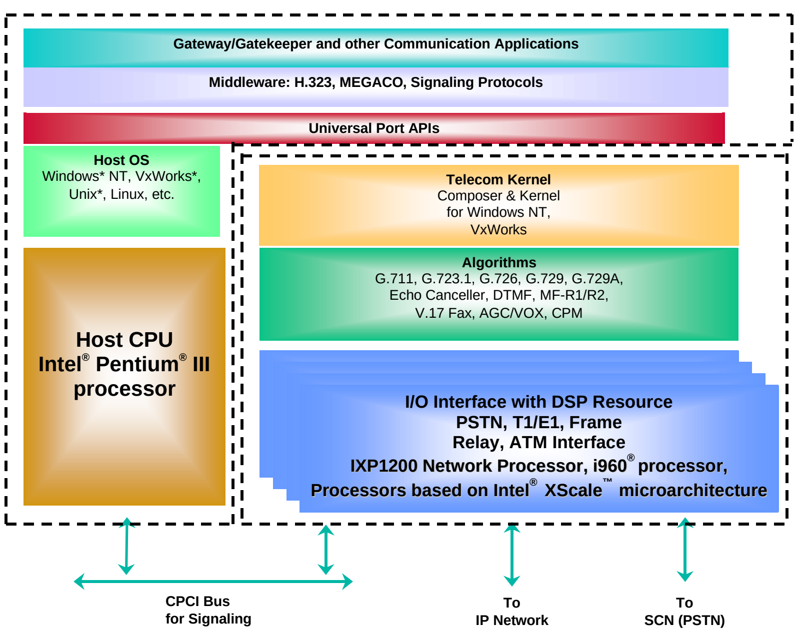

Figure 1 - Intel® Internet Exchange Architecture Gateway Implementation

density, speed, and complexity of I/O functions, these boards can utilize the Intel® IXP1200 network processor, processors based on the Intel® XScale™ microarchitecture or Intel® i960® processors.

Voice over IP: General Considerations

Before Voice over IP permeates the network and becomes the leading technology of choice for voice applications, perceived voice quality must be made comparable to the quality available from traditional wire line telephony provided by PSTN carriers. IP technology has certain advantages for carrying data but has limits in voice or other isochronous services. The best-effort procedure adopted in IP is good for store-and-forward type communications, but it may not be able to provide adequate quality for voice. This is particularly true in public networks because of prioritization and loading considerations. Delay can create quality issues in bi-directional communications applications, while jitter degrades voice clarity. The IETF (Internet Engineering Task Force) and ITU (International Telecommunication Union) are working to enable the IP network to address the delay and jitter requirements of voice.

Co-existence of VoIP and PSTN

The evolution of IP telephony and its integration into everyday operations will require the evolution of the underlying infrastructure. Until public IP networks entirely replace the major elements of traditional telecommunication networks, the IP network and the PSTN will need to coexist. This is why Inter Working Functions (IWF) are currently defined, and why gateways/gatekeepers are installed to bridge between the two networks. Dissimilarities in signaling, control, and media translation are being absorbed in IWF and gateways to support end-to-end connectivity.

A gateway can be a single piece of equipment handling both media and control parts of a connection between the two networks. This is known as a composed gateway. There are also several Gateway Control Protocols (GCP) that will allow these functions to be spread over multiple pieces of specialized equipment. Such a "decomposed gateway" typically consists of a signaling gateway with an external or internal call agent and one or more media gateways, with some form of communication provided between these elements. Depending on the application scenario (port density, service provider, and so on), a gateway can be either composed or decomposed.

Signaling in PSTN and IP

Signaling in the telecommunications world refers to information associated with call set-up, call progress, routing, and teardown across a physical or virtual circuit. The signaling protocol in modern circuit-switched networks is based on Signaling System 7 (SS7). SS7 is a dedicated network for call control and it is separate from the transmission facility dealing with the actual voice connection.

Unlike voice telephony, the IP network requires extended signaling due to the different nature of the end-points. The end-points can support a multitude of capabilities with respect to bandwidth requirements, audio, video connection, and data capabilities. Before the end-points can establish the session and engage in conversation, it is necessary to ensure that they share the same capabilities.

VoIP Protocols

Call control is a type of signaling related to the connection of a call between two parties. When the two parties are attached to different networks, the gateway between the networks will ascertain the translation of signaling and media type conversion.

Figure 2 illustrates the architectural concept where the user application interfaces with the IP-telephony signaling stack that uses IP as an end-to-end protocol. The user application can be a vendor-specific voice gateway implementation. The IP Telephony (IPT) stack handles the communication between the two applications and the control entity.

There are a couple of initiatives from international forums that address IPT protocol requirements. These include Recommendation H.323 from ITU and Session Initiation Protocol (SIP-RFC2543) from IETF. The following sections describe in detail the H.323 protocol, its architecture, and typical applications in a gateway implementation.

H.323 Protocol and Revisions

H.323 is an ITU recommendation and is the most widely used and best proven protocol in the VoIP arena. H.323 is a family of standards that specify the components, protocol, and procedures required to provide multimedia communications over packet-based networks, including those based on IP. H.323 supports point-to-point and multi-point communications using four types of components. These components include Terminals, Gateways, Gatekeepers, and Multi-point Conferencing Units (MCU). By complying with H.323, multimedia products and applications from multiple vendors can interoperate, enabling users to communicate without the concern for compatibility.

The ITU's Study Group 16 approved the H.323 specification in 1996. Version 2 was approved in January 1998. The standard is broad in scope and includes both stand-alone devices and embedded personal computer technology in addition to point-to-point and multipoint conferences.

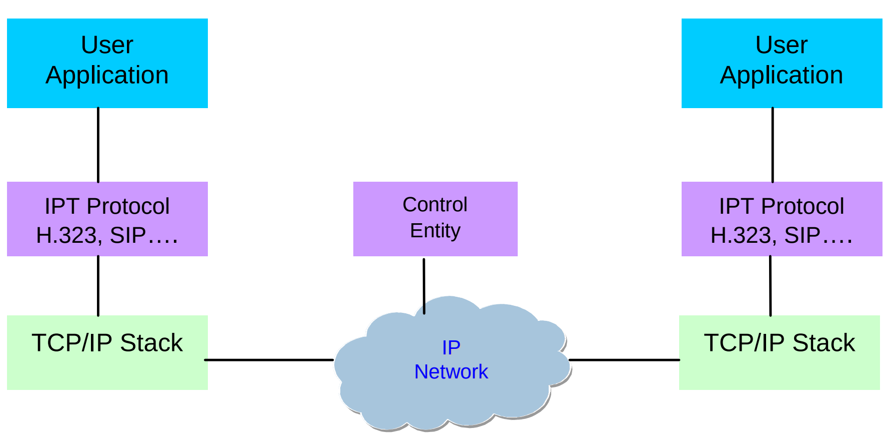

Figure 2 - IP Telephony Signaling Architecture

H.323 also addresses call control, multimedia management, and bandwidth management as well as interfaces between LANs and other networks. Version 2 of the H.323 standard addressed deficiencies in Version 1 and introduced new functionality within the existing protocols, such as Q.931, H.245 and H.225. The most significant advances were in security, fast call setup and supplementary services.

Bridging Between PSTN and IP via the H.323 Gateway

A gateway provides translation of protocols for call set-up, release, and conversion of media formats between different networks. An application of the H.323 gateway is IP telephony, where the H.323 gateway interconnects an IP network and switched circuit network (SCN) such as ISDN. For example, a gateway can connect and provide communication between an H.323 terminal and SCN networks (SCN networks include all switched telephony networks including PSTN).

This connectivity of dissimilar networks is achieved through the use of translation protocols for call set-up and release, converting media formats between different networks, and then transferring information between the networks connected to the gateway.

On the H.323 side, as shown in Figure 3, the gateway runs H.245 control signaling for exchange capabilities, H.225 call signaling for call set-up and release, and H.225 registration, admissions, and status (RAS) for registration with the gatekeeper. On the SCN side, a gateway runs SCN-specific protocols (such as ISDN and

SS7 protocols). Terminals communicate with gateways using the H.245 control-signaling protocol and H.225 call-signaling protocol.

The gateway translates these protocols in a transparent fashion to the respective counterparts on the non-H.323 network and vice versa. The gateway also performs call set-up and clearing on both the H.323 network side and the non-H.323 network side. Translation between audio, video, and data formats may also be performed by the gateway. Audio and video translation may not be required if both terminal types find a common communications mode. For example, in the case of a gateway to H.320 terminals on the ISDN, both terminal types require G.711 audio and H.261 video, so a common mode always exists.

Building a VoIP Gateway

Equipment vendors can cost effectively build the gateway function by embedding the VoIP technology within their legacy telecom or datacom equipment. The basic concept of VoIP is to enable the telephones and fax machines to work over the IP network or Internet. Gateways are very important components because they provide a way for users of legacy telephone equipment to realize the benefits of packet switching. With the massive installed base of traditional telephone equipment, gateways present a significant opportunity to telecom and datacom equipment manufacturers.

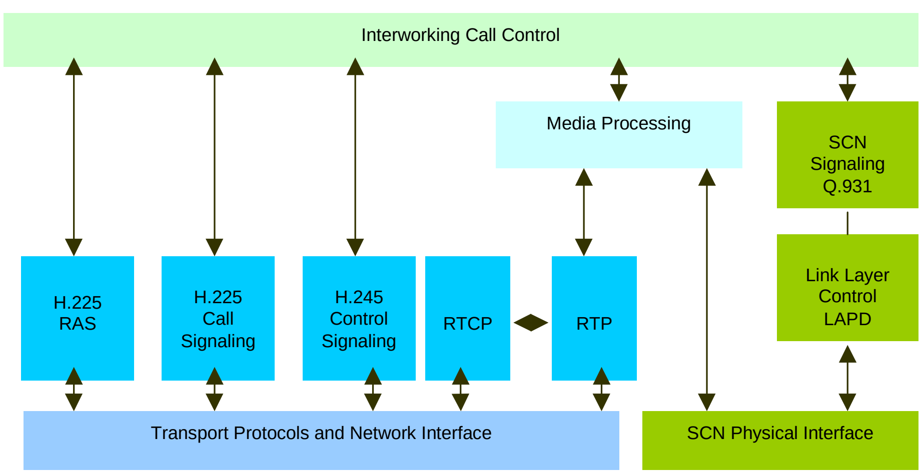

Figure 3 - Gateway Functions Based on H.323

PC vs. Embedded Solutions

A quick and low cost solution to achieve the gateway functionality would be to use PC-based plug-in cards with telephony interfaces and DSP resources. In fact, this has been the trend in the early days of VoIP telephony products. However, the acceptance of these solutions is low, due to their voice quality and limited scalability.

An embedded product is built specifically for the function it provides, with dedicated hardware, a suitable operating system and an appropriate hardware/software architecture designed for the application. Although built from off-the-shelf components, these solutions are helped by several industry initiatives like cPCI (Compact PCI), H.100/110, hot swap, IPMI (Intelligent Platform Management Interface), and Web-based management, to meet the higher reliability and redundancy requirements of today's telephone infrastructure.

One advantage of cPCI for the board and system designer is that the same electrical PCI bus used in PCs is employed as an on-card bus and as a system bus for cPCI cards. This allows the designers to take advantage of the chipsets and the new technologies available for PC applications. Since many Computer Telephony Integration (CTI) applications are largely PC-based today, they can easily be migrated to the cPCI form factor. In addition, an embedded gateway solution has superior voice quality, high reliability, lower cost, high density and scalability, compared to PC-based versions.

PC-based platforms typically involve lower up-front costs and initial engineering investment compared to embedded solutions. However recurring per-channel costs are much lower for an embedded solution in production volumes. In addition, higher port density and scalability can be realized in an embedded solution, which makes it the preferred choice for carrier-class applications.

Hardware Recommendations

Whether a developer is building a new product or adding IP capability to existing voice equipment, new hardware is often required. A functional block diagram of the components necessary to build a gateway between the PSTN and IP network is shown in Figure 4.

The processing requirements of a VoIP solution are split between the digital signal processor (DSP) and the microprocessor. DSPs feature a specialized architecture that makes them suitable for the repetitive calculations required for voice coding, compression, echo cancellation, fax modem, and other real-time voice specific implementations. Intel® Pentium® processors with MMX™ technology and Single Instruction Multiple Data (SIMD) capabilities can be used to implement these signal processing functions if the port density is adequate for the available MIPS, thus minimizing or eliminating the need for separate DSPs. In the case of carrier-class VoIP products, where a large number of channels are being handled, it is better to handle media related functions with dedicated DSP resources, leaving signaling, control and management functions to the host processor.

The selection of a DSP is based on several criteria, port density being an important factor. The number of voice and fax channels that can be served by a DSP depends on the available MIPS and the memory to run the algorithms.

The choice of microprocessor depends on the processing power required for the application, in addition to scalability requirements and cost. The selected processor needs to handle signaling, and management and control functions as required by the application. In the case of the carrier-class VoIP gateway, the microprocessor is required to handle the control portion of H.323. Program and data memory is also required in the system for the microprocessor and the DSP. Non-volatile memory is needed for program storage.

As shown in Figure 4, the connection to the PSTN can be through digital interfaces like T1 or E1, depending on the geographical location. The T1 or E1 interface multiplexes 24 (T1) or 30 (E1) channels of voice, along with signaling, into a single data stream. The connection to PSTN can also be made through an analog interface if there is an interposed codec to convert the analog signals into an appropriately coded digital stream. The received voice channels are fed to the DSP for handling the media portion, while the signaling portion is connected to the microprocessor.

Software Requirements

Software requirements can be considered in two parts. The first part is related to packetization of voice from PSTN, placing it on the data network, and vice versa. This requires a core software functionality for voice compression/decompression, echo cancellation, tone processing and jitter removal. The other part of the software deals with the telephony signaling such as on-hook/off-hook recognition and converting them to connect/disconnect messages that can be understood in the IP environment.

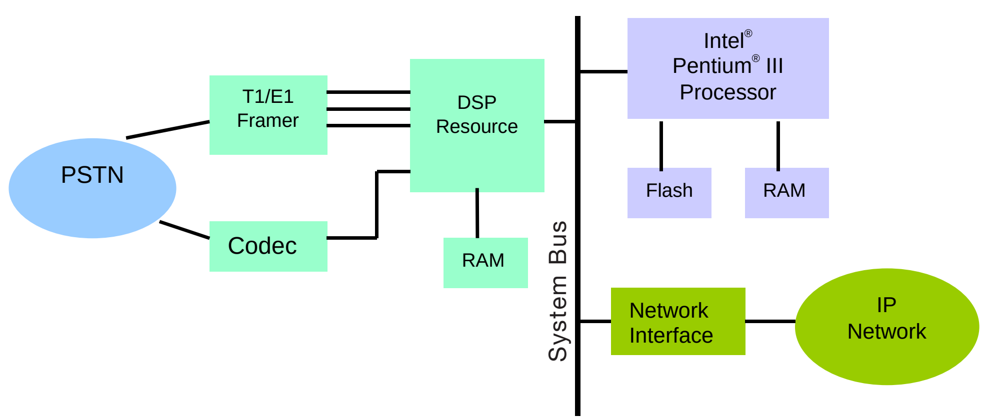

Figure 4 - Hardware Functions of PSTN to IP Gateway

The media processing block, which contains signal processing, coding, compression, and echo cancellation functions, in addition to packetization of voice and real-time protocols, handles the interface to PSTN and voice processing. The media processing block also generates the call progress tones to the user and DTMF digits towards the network. Echo cancellation is an essential part of the media processing block, particularly in VoIP networks, due to the increased latency associated with voice compression algorithms.

Call processing essentially consists of dealing with telephony signaling and translation into H.323 signaling that can be understood by the IP network. The TCP/UDP and IP stack, with suitable layer 2 and physical layer functionality, handle the network interface to the IP network.

Figure 5 illustrates embedded software requirements for a VoIP gateway.

System Architecture

Figure 6 illustrates the system architecture of an H.323 gateway.

The signaling function is implemented in the host CPU, and the protocols relating to media traffic are accomplished in the I/O plane. H.323 signaling elements, including H.225 call signaling, H.245 control, and H.245 RAS, are implemented in the host CPU. The host CPU also handles SCN call signaling, Q.931. The media protocols for VoIP network (RTP/RTCP, transport layer TCP/IP) and the link layer protocol for SCN (LAPD) are mapped on the I/O plane.

The media processing module shown in Figure 6 performs media streaming operations including voice compression, voice detection, gain control, jitter compensation and echo cancellation, while switching the voice traffic between the IP network and SCN. The media processing module must support G.711 for basic compression and may also support other compression algorithms to minimize the network bandwidth requirement. Voice compression and the echo cancellation may be implemented in software or with a special-purpose DSP chip. The choice depends on the port density and the target cost of equipment. The usage of a dedicated DSP chip may decrease processing time, and hence latency, but can increase the overall cost of the equipment. On the other hand, software implementation of media protocols on a host processor requires careful analysis of available MIPS in the host and capability to support DSP-like signal processing functions.

The Ethernet driver and the serial driver are the respective driver module for the Ethernet controller and the Serial controller.

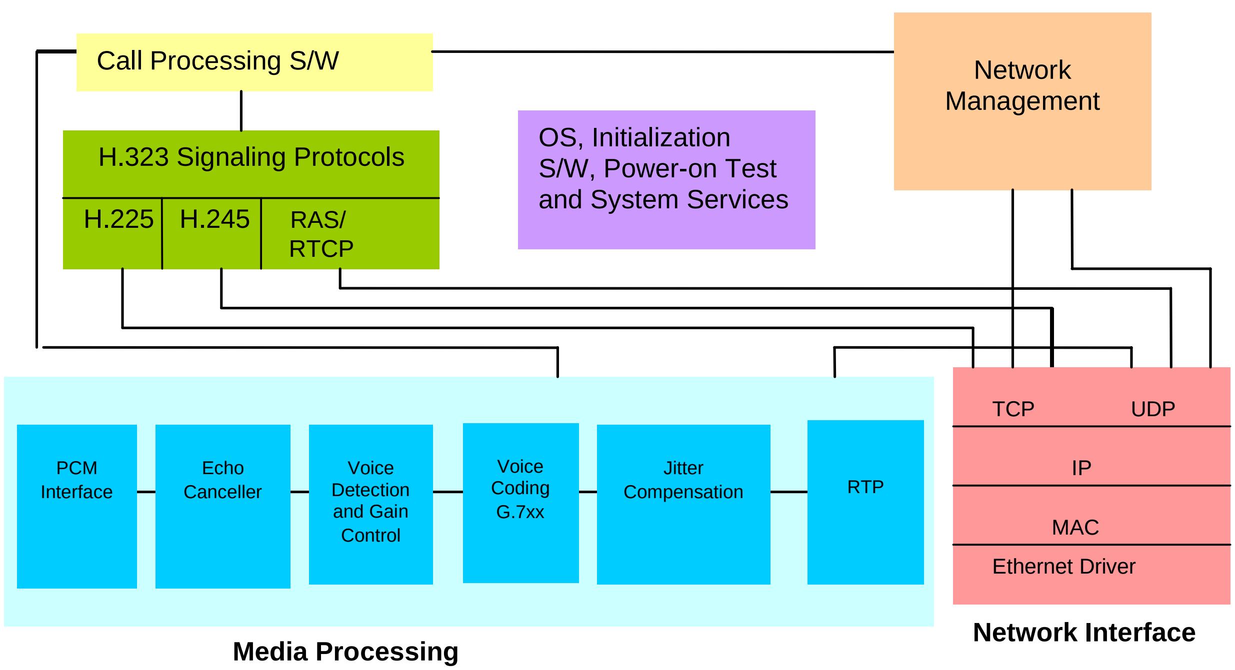

Figure 5 - Software Functions in H.323 Gateway

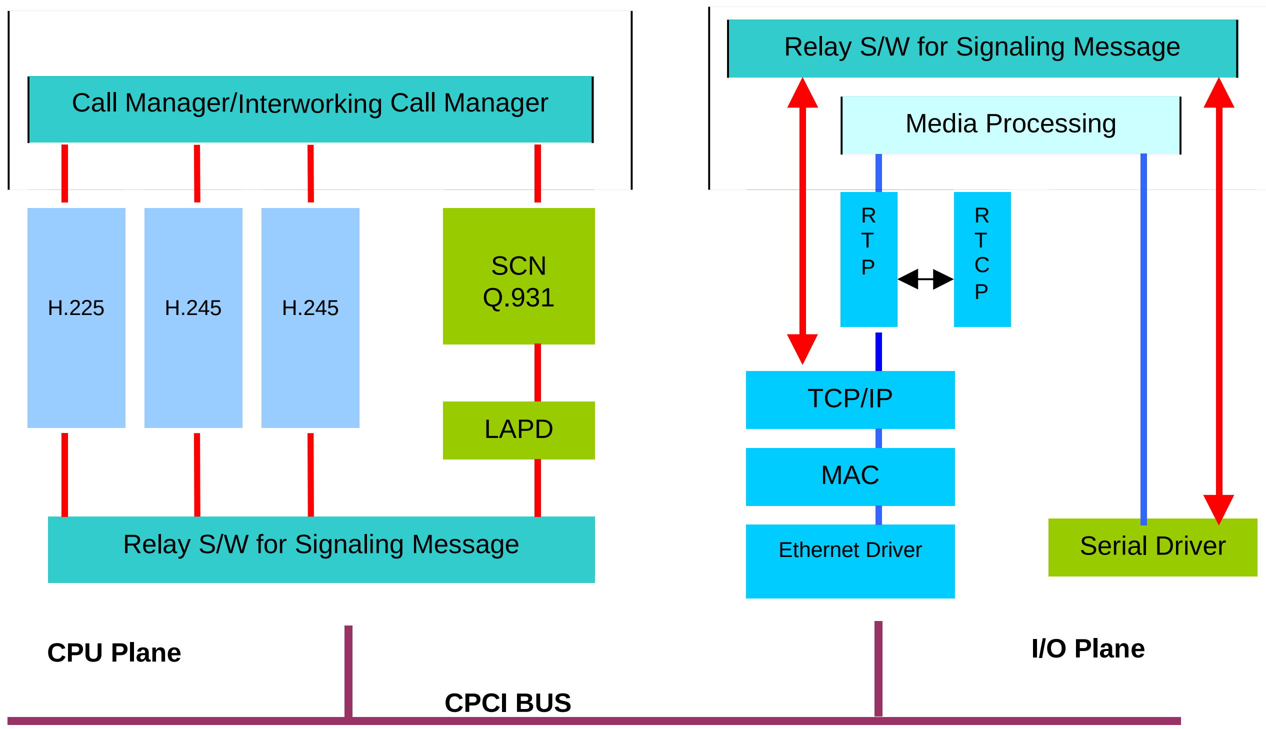

Figure 6 - System Implementation of an H.323 Gateway

The system shown in Figure 6 switches the media traffic between the two networks (IP and SCN) in the I/O plane without loading the PCI BUS. The signaling messages, both in the IP network and the SCN network, are delivered to the signaling modules in the host CPU. A separate software relay module, as shown in the diagram, passes the signaling messages between the CPU and the I/O plane over the PCI BUS. At the end of the signaling phase, the media channels are opened between the SCN and IP network. The switching of media channels between the SCN and IP network is done entirely in the I/O plane, ensuring that the PCI bus will not be a bottleneck.

PCI Bandwidth Considerations for H.323 Signaling

The bandwidth requirement for H.323 signaling can be calculated based on the number of calls and the typical size of each signaling message exchanged over the PCI BUS between the I/O processor plane and the host CPU. The PCI bus carries only the signaling messages for the two networks, and it is assumed that all the media related switching occurs in the I/O plane.

Bandwidth considerations include the following:

1. Message size for H.323 signaling
   - Total size of all messages exchanged for call establishment is approximately 600 bytes
   - Total size of messages exchanged for call release is approximately 53 bytes
   - Total size of RAS messages exchanged is approximately 140 bytes.
2. Message size for SCN signaling
   - Total size of all messages exchanged for connection establishment is approximately 70 bytes
   - Total size of messages exchanged for call release is approximately 20 bytes.

Based on the above message sizes, the total number of message bytes exchanged between the I/O plane and the CPU for establishing and releasing a call across the SCN and the IP network, and typical number of calls that the PCI bus can handle, is calculated below:

The total number of message bytes per call is approximately 900 bytes. The total PCI bandwidth required for establishing and releasing 1000 calls per sec is 0.9 Mbytes/sec.

Assuming the PCI bandwidth at 32 bit and 33 MHz is 132 Mbytes/sec, and an effective rate of 70 percent, the number of calls that can be handled by the PCI bus is 100,000.

As these calculations demonstrate, the PCI bus has enough bandwidth to handle signaling transactions for 100,000 calls between the two networks. In a typical scenario where one processor is used to handle signaling for SCN and IP networks, it is obvious that a PCI bus can be used to carry signaling information from several I/O peripheral cards to the host processor. This allows enough scalability in building medium to large carrier-class gateways supporting a few hundred to thousands of ports.

Applied Computing Platform Providers and Software Suppliers

Time-to-market is a key factor in the development of communications equipment, particularly products dealing with the Internet infrastructure. With this view in mind, The Intel® Applied Computing Platform Providers (ACPP) program has been established to support the development of embedded Intel Architecture platforms for communications applications. In addition, Intel is working with vendors to port, optimize and validate new signaling, media control and management stacks on Intel processors.

Conclusion

A carrier-class VoIP gateway can be built around Intel Internet Exchange Architecture by optimally partitioning the control and signaling layers and the media protocol layers between the host CPU and the I/O processors. Intel provides a comprehensive range of processors and associated chipsets for both the host and I/O peripheral boards. In addition, Intel has launched key programs with platform and software vendors to accelerate the development of communication platforms and gateways that enable new and enhanced services for the Internet.

For More Information

For more information on Intel's Applied Computing Platform Provider Program, visit developer.intel.com/platforms/applied/acpp/index.htm

For more information on telephony software stacks, visit www.trillium.com

For more information on Intel Internet Exchange Architecture solutions, visit www.intel.com/ixa

Order Number: 273450-001 Printed in USA/1000/7.5K/IL2110L LY